



Microsoft

70-463 Exam

Microsoft Implementing a Data Warehouse with Microsoft SQL Server 2012/2014 Exam

Thank you for Downloading 70-463 exam PDF Demo

You can Buy Latest 70-463 Full Version Download

<https://www.certkillers.net/Exam/70-463>

<https://www.certkillers.net>

Version: 18.0

Question: 1

You are developing a project that contains multiple SQL Server Integration Services (SSIS) packages. The packages will be deployed to the SSIS catalog. One of the steps in each package accesses an FTP site to download sales transaction data.

You create project parameters to store the username and password that are used to access the FTP site.

You need to ensure that the username and password values are encrypted when they are deployed. What should you do?

- A. Set the Sensitive property of the parameters to True.
- B. Set the ProtectionLevel property of the package to EncryptSensitiveWithUserKey.
- C. Change the parameters to package parameters.
- D. Change the project to the Legacy Deployment model.

Answer: A

Question: 2

You develop a SQL Server Integration Services (SSIS) package that imports SQL Azure data into a data warehouse every night.

The SQL Azure data contains many misspellings and variations of abbreviations. To import the data, a developer used the Fuzzy Lookup transformation to choose the closest-matching string from a reference table of allowed values. The number of rows in the reference table is very large.

If no acceptable match is found, the Fuzzy Lookup transformation passes a null value.

The current setting for the Fuzzy Lookup similarity threshold is 0.50.

Many values are incorrectly matched.

You need to ensure that more accurate matches are made by the Fuzzy Lookup transformation without degrading performance.

What should you do?

- A. Change the Exhaustive property to True.
- B. Change the similarity threshold to 0.55.
- C. Change the similarity threshold to 0.40.
- D. Increase the maximum number of matches per lookup.

Answer: B

Explanation:

<http://msdn.microsoft.com/en-us/library/ms137786.aspx>

Question: 3

You install a SQL Server 2012 database engine instance on a production server. A month later, you

install SQL Server 2012 Integration Services (SSIS).

You must develop an SSIS project and deploy it to the server by using the Project Deployment model. Operations Log records that are outside the configured retention period must be cleaned automatically.

You need to create the SSIS catalog on the production server.

What should you do? (Each correct answer presents part of the solution. Choose all that apply.)

- A. Enable XP Command Shell.
- B. Enable CLR Integration.
- C. Enable OLE Automation.
- D. Start the SQL Server Browser service.
- E. Enable Cross Database Ownership Chaining
- F. Start the SQL Server Agent service
- G. Enable Ad Hoc Remote Queries.

Answer: B, F

Ref: [http://msdn.microsoft.com/en-us/library/gg471509\(v=sql.110\).aspx](http://msdn.microsoft.com/en-us/library/gg471509(v=sql.110).aspx)

Question: 4

DRAG DROP

A SQL Server Integration Services (SSIS) package named DataFeed interacts with an external vendor data feed. The package is executed several times a day, either as part of other packages' control flow or by itself. The external data feed is unreliable because network failures and slow response times are frequent. The package is currently deployed on the file system.

To analyze the reliability of the external data feed, you must collect execution data.

Every time the DataFeed package is executed, the following information must be logged:

Start Time

End Time

Execution Result

Execution Duration

You need to design a logging solution that meets the requirements by using the least amount of administrative and development effort.

Which three actions should you perform in sequence? (To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.)

Add OnError and OnWarning event handlers.	
Query the catalog.executable_statistics view.	
Query the ExecutionLog table.	
Deploy the project that contains the package to the SSIS catalog and execute the package on the server.	
Deploy the package to the msdb database.	
Add an Execute SQL task to the event handlers.	
Query the catalog.execution_data_statistics view.	
Create a new project and add the package to the project.	

Answer:

Box 1:

Create a new project and add the package to the project.

Box 2:

Deploy the project that contains the package to the SSIS catalog and execute the package on the server.

Box 3:

Query the catalog.executable_statistics view.

Ref: <http://msdn.microsoft.com/en-us/library/hh479592.aspx>

Question: 5

A SQL Server Integration Services (SSIS) package on a computer is failing periodically in production. The package was created less than one year ago and was deployed to the SSIS catalog. Sometimes the package is started on a SQL Agent schedule; sometimes the package is started manually by an SSIS developer by using the Object Explorer in SQL Server Management Studio. You need to identify the authenticated user responsible for starting the package each time it failed in the past.

Where can you find this information?

- A. the SQL Server Log
- B. the SSISDB.[catalog].[executions] view
- C. the SSISDB.[catalog].[event_messages] view
- D. the SQL Agent Job History
- E. the SQL Agent Error Log

Answer: B

Question: 6

You maintain a SQL Server Integration Services (SSIS) package. The package was developed by using SQL Server 2008 Business Intelligence Development Studio (BIDS).

The package includes custom scripts that must be upgraded.

You need to upgrade the package to SQL Server 2012.

Which tool should you use?

- A. SSIS Upgrade Wizard in SQL Server 2008 BIDS
- B. SSIS Upgrade Wizard in SQL Server Data Tools
- C. SQL Server DTExecUI utility (dtexecui.exe)
- D. SQL Server dtexec utility (dtexec.exe)

Answer: B

Question: 7

DRAG DROP

You are editing a SQL Server Integration Services (SSIS) package that contains a task with a sensitive

property.

You need to create a project parameter and configure it so that its value is encrypted when it is deployed to the SSIS catalog.

Which three steps should you perform in sequence? (To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.)

Select the property to be parameterized and set the scope to **Project**.

Edit the parameter and set the **Sensitive** property to **True**.

Convert the project to the Legacy Deployment model.

Select the property to be parameterized and set the scope to **Package**.

Right-click the task and choose **Parameterize**.

Right-click the task and choose **Properties**.



Answer:

Select the property to be parameterized and set the scope to **Project**.

Right-click the task and choose **Parameterize**.

Convert the project to the Legacy Deployment model.

Select the property to be parameterized and set the scope to **Package**.

Edit the parameter and set the **Sensitive** property to **True**.

Right-click the task and choose **Properties**.

Question: 8

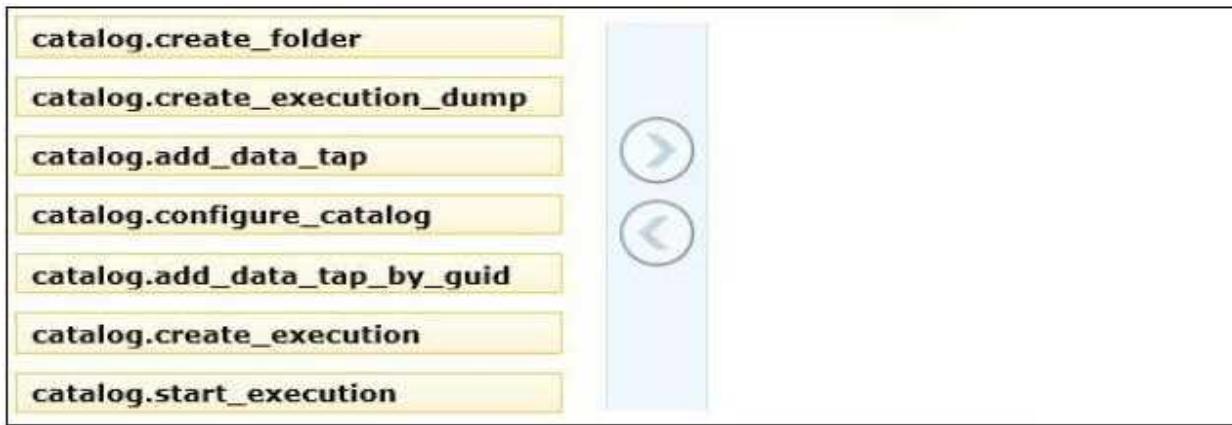
DRAG DROP

A new SQL Server Integration Services (SSIS) project is deployed to the SSIS catalog.

To troubleshoot some data issues, you must output the data streaming through several data flows into text files for further analysis. You have the list of data flow package paths and identification strings of the various task components that must be analyzed.

You need to create these output files with the least amount of administrative and development effort.

Which three stored procedures should you execute in sequence? (To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.)



Answer:

Box 1:

catalog.create_execution

Box 2:

catalog.add_data_tap

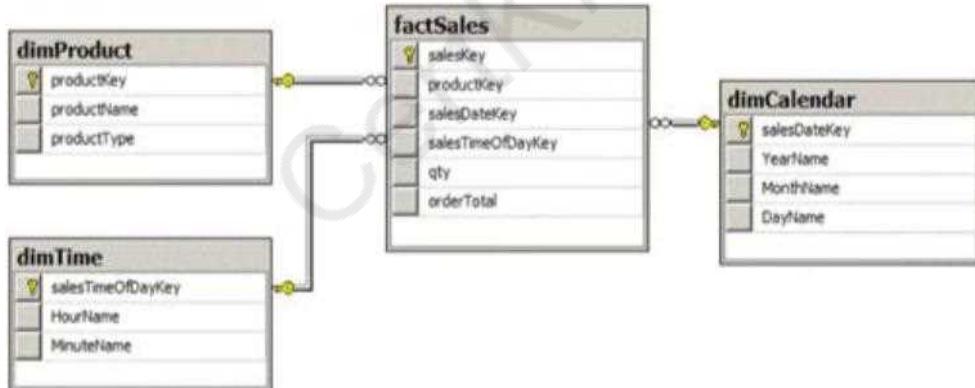
Box 3:

catalog.start_execution

Ref: [http://msdn.microsoft.com/en-gb/library/hh230989\(v=sql.110\).aspx](http://msdn.microsoft.com/en-gb/library/hh230989(v=sql.110).aspx)

Question: 9

You are reviewing the design of an existing fact table named factSales, which is loaded from a SQL Azure database by a SQL Server Integration Services (SSIS) package each day. The fact table has approximately 1 billion rows and is dimensioned by product, sales date, and sales time of day.



The database administrator is concerned about the growth of the database. Users report poor reporting performance against this database. Reporting requirements have recently changed and the only remaining report that uses this fact table reports sales by product name, sale month, and sale year. No other reports will be created against this table.

You need to reduce the report processing time and minimize the growth of the database.

What should you do?

- A. Partition the table by product type.
- B. Create a view over the fact table to aggregate sales by month.
- C. Change the granularity of the fact table to month.

D. Create an indexed view over the fact table to aggregate sales by month.

Answer: C

Question: 10

You are designing a data warehouse with two fact tables. The first table contains sales per month and the second table contains orders per day.

Referential integrity must be enforced declaratively.

You need to design a solution that can join a single time dimension to both fact tables.

What should you do?

- A. Create a time mapping table.
- B. Change the level of granularity in both fact tables to be the same.
- C. Merge the fact tables.
- D. Create a view on the sales table.

Answer: B

Question: 11

You are designing a data warehouse for a software distribution business that stores sales by software title. It stores sales targets by software category. Software titles are classified into subcategories and categories. Each software title is included in only a single software subcategory, and each subcategory is included in only a single category. The data warehouse will be a data source for an Analysis Services cube.

The data warehouse contains two fact tables:

factSales, used to record daily sales by software title

factTarget, used to record the monthly sales targets by software category

Reports must be developed against the warehouse that reports sales by software title, category and subcategory, and sales targets.

You need to design the software title dimension. The solution should use as few tables as possible while supporting all the requirements.

What should you do?

- A. Create three software tables, dimSoftware, dimSoftwareCategory, and dimSoftwareSubcategory and a fourth bridge table that joins software titles to their appropriate category and subcategory table records with foreign key constraints. Direct the cube developer to use key granularity attributes.
- B. Create three software tables, dimSoftware, dimSoftwareCategory, and dimSoftwareSubcategory. Connect factSales to all three tables and connect factTarget to dimSoftwareCategory with foreign key constraints. Direct the cube developer to use key granularity attributes.
- C. Create one table, dimSoftware, which contains Software Detail, Category, and Subcategory columns. Connect factSales to dimSoftware with a foreign key constraint. Direct the cube developer to use a non-key granularity attribute for factTarget.
- D. Create two tables, dimSoftware and dimSoftwareCategory. Connect factSales to dimSoftware and factTarget to dimSoftwareCategory with foreign key constraints. Direct the cube developer to use key

granularity attributes.

Answer: C

Question: 12

You are designing a data warehouse hosted on SQL Azure. The data warehouse currently includes the dimUser and dimDistrict dimension tables and the factSales fact table. The dimUser table contains records for each user permitted to run reports against the warehouse; and the dimDistrict table contains information about sales districts.

The system is accessed by users from certain districts, as well as by area supervisors and users from the corporate headquarters.

You need to design a table structure to ensure that certain users can see sales data for only certain districts. Some users must be permitted to see sales data from multiple districts.

What should you do?

- A. Add a district column to the dimUser table.
- B. Partition the factSales table on the district column.
- C. Create a userDistrict table that contains primary key columns from the dimUser and dimDistrict tables.
- D. For each district, create a view of the factSales table that includes a WHERE clause for the district.

Answer: C

Question: 13

You are reviewing the design of a customer dimension table in an existing data warehouse hosted on SQL Azure.

The current dimension design does not allow the retention of historical changes to customer attributes such as Postcode.

You need to redesign the dimension to enable the full historical reporting of changes to multiple customer attributes including Postcode.

What should you do?

- A. Add StartDate and EndDate columns to the customer dimension.
- B. Add an IsCurrent column to the customer dimension.
- C. Enable Snapshot Isolation on the data warehouse.
- D. Add CurrentValue and PreviousValue columns to the customer dimension.

Answer: A

Question: 14

You are implementing the indexing strategy for a fact table in a data warehouse. The fact table is named Quotes. The table has no indexes and consists of seven columns:

- [ID]

- [QuoteDate]
- [Open]
- [Close]
- [High]
- [Low]
- [Volume]

Each of the following queries must be able to use a columnstore index:

- `SELECT AVG ([Close]) AS [AverageClose] FROM Quotes WHERE [QuoteDate] BETWEEN '20100101' AND '20101231'.`
- `SELECT AVG([High] - [Low]) AS [AverageRange] FROM Quotes WHERE [QuoteDate] BETWEEN '20100101' AND '20101231'.`
- `SELECT SUM([Volume]) AS [SumVolume] FROM Quotes WHERE [QuoteDate] BETWEEN '20100101' AND '20101231'.`

You need to ensure that the indexing strategy meets the requirements. The strategy must also minimize the number and size of the indexes.

What should you do?

A. Create one columnstore index that contains [ID], [Close], [High], [Low], [Volume], and [QuoteDate].

B. Create three columnstore indexes:

One containing [QuoteDate] and [Close]

One containing [QuoteDate], [High], and [Low]

One containing [QuoteDate] and [Volume]

C. Create one columnstore index that contains [QuoteDate], [Close], [High], [Low], and [Volume].

D. Create two columnstore indexes:

One containing [ID], [QuoteDate], [Volume], and [Close]

One containing [ID], [QuoteDate], [High], and [Low]

Answer: C

Explanation:

Reference: <http://msdn.microsoft.com/en-us/library/gg492088.aspx>

Reference: <http://msdn.microsoft.com/en-us/library/gg492153.aspx>

Question: 15

You are designing an enterprise star schema that will consolidate data from three independent data marts. One of the data marts is hosted on SQL Azure.

Most of the dimensions have the same structure and content. However, the geography dimension is slightly different in each data mart.

You need to design a consolidated dimensional structure that will be easy to maintain while ensuring that all dimensional data from the three original solutions is represented.

What should you do?

A. Create a junk dimension for the geography dimension.

B. Implement change data capture.

C. Create a conformed dimension for the geography dimension.

D. Create three geography dimensions.

Answer: C

CertKillers.net

Thank You for trying 70-463 PDF Demo

To Buy Latest 70-463 Full Version Download visit link below

<https://www.certkillers.net/Exam/70-463>

Start Your 70-463 Preparation

[Limited Time Offer] Use Coupon “**CKNET**” for Further discount on your purchase. Test your 70-463 preparation with actual exam questions.

<https://www.certkillers.net>